arXiv:1602.06589v3 [cs.CE] 15 Aug 2016

High-performance generation of the Hamiltonian and Overlap matrices in FLAPW methods

Edoardo Di Napoli[a,b,*], Elmar Peise[b], Markus Hrywniak[c], Paolo Bientinesi[b]

[a] *Jülich Supercomputing Centre, Forschungszentrum Jülich, and JARA, 52425 Jülich, Germany.*
[b] *RWTH Aachen University, Aachen Institute for Advanced Study in Computational Engineering Science, Schinkelstr. 2, 52062 Aachen, Germany*
[c] *RWTH Aachen University, German Research School for Simulation Sciences, Schinkelstr. 2a, 52062 Aachen, Germany*

Abstract

One of the greatest efforts of computational scientists is to translate the mathematical model describing a class of physical phenomena into large and complex codes. Many of these codes face the difficulty of implementing the mathematical operations in the model in terms of low level optimized kernels offering both performance and portability. Legacy codes suffer from the additional curse of rigid design choices based on outdated performance metrics (e.g. minimization of memory footprint). Using a representative code from the Materials Science community, we propose a methodology to restructure the most expensive operations in terms of an optimized combination of dense linear algebra (BLAS3) kernels. The resulting algorithm guarantees an increased performance and an extended life span of this code, enabling larger scale simulations.

Keywords: Density Functional Theory, high-performance computing, dense linear algebra, matrix generation, performance portability, FLAPW, FLEUR

1. Introduction

In this paper, we look at the issues of performance portability and extensibility of legacy codes in scientific computing. As a case study, we consider FLEUR [1], a code for electronic structure calculations. FLEUR was developed at the Forschungszentrum Jülich for Materials Science simulations over the course of 2 decades. As such, it has grown into an extensive project with numerous features, spread over more than 100 000 lines of code. However, as is the case for many legacy codes, its incremental and functionality-oriented design resulted in an application with poor use of modern hardware capabilities. Unfortunately, modifying the existing code to exploit parallelism and to increase its modularity so as to allow for the use of external high-performance libraries has proven quite vexing. Therefore, we follow a clean-slate approach: starting from the mathematical description of a major portion of FLEUR, we develop a modular algorithm that employs high-level linear algebra operations implemented

*Principal corresponding author
Email addresses: e.di.napoli@fz-juelich.de (Edoardo Di Napoli), peise@aices.rwth-aachen.de (Elmar Peise), markus.hrywniak@rwth-aachen.de (Markus Hrywniak), pauldj@aices.rwth-aachen.de (Paolo Bientinesi)

in optimized libraries. The resulting implementation distinctly outperforms the original one and, thanks to its modularity, guarantees excellent performance portability. While requiring a larger initial effort compared to the traditional approach of incremental parallelization and optimization, such a complete redevelopment of legacy software yields, in the long run, remarkable benefits in terms of performance, portability, and maintainability.

Large legacy codes, such as FLEUR, are very common in computational science. Well-tested and with validated results, most simulation and experimental codes have a rich life span lasting years and even decades. The initial implementation usually resembles closely the mathematical formulation of the physical problem, and often it is the direct translation of such formulation into code. By abstracting from specific hardware, this approach allows for fast result validation; indeed, it is so natural that even inspired the name of "Fortran" (*For*mula *Tran*slation), the oldest programming language in use, and still one of the most widespread. Despite its simplistic nature, this type of code development has been extremely successful, enabling great scientific breakthroughs that in some cases were even awarded with the Nobel prize[1].

In spite of their initial successes, these legacy codes were often implemented without keeping in mind the necessity of a layered structure which would allow for the extension of the code by adding new features. Nor did those designing the code forecast the necessity of running larger, more complex, and more accurate simulations which would require enhanced parallelism. In the same fashion, no systematic approach was undertaken for a carefully engineered exploitation of processors' architectural features in order to avoid computational bottlenecks. For instance, as more functionalities were added to an early implementation, more and more "premature optimizations" [22] made their way into the codebase. Such practices have the profound consequence of making the job of exposing parallelism very onerous if not punishing. Consequently, in order to enhance their portability to massively parallel supercomputers, legacy codes have to undergo a substantial restructuring (see for example [**? ?**]).

At the time of their inception, most legacy codes had to deal with memory limitations. Not only was computer memory expensive, and thus limited in size, but also, Fortran did not allow for dynamical allocation of memory, forcing programmers to quantify in advance the exact size of the allocated working space in the physical memory. Most codes relied on a reduction of performed floating point operations (FLOP) for speedup. The introduction of a hierarchy of caches changed the paradigm, however this hardware revolution was hardly noted by most of the community of computational scientists. The by-product of the change in hardware architecture produced a significant paradigm shift: although memory usage and FLOP count are still valid metrics, they are not synonymous with computational efficiency.

On modern computing architectures with cache hierarchies, "unqualified" FLOP count may lead to drawing incorrect conclusions over which are the optimal algorithmic choices. For instance, algorithms that perform the exact same number of FLOPs to execute a low level operation can easily be one order of magnitude apart in terms of execution time. Consequently, using the minimization of FLOPs as a base for algorithm choice does not necessarily imply a lower execution time [2]. A natural corollary to such a statement is that – if necessary – one can trade away a lower number of "slow" FLOPs for a larger number of "fast" ones [3, 4].

[1]The 1998 Nobel Prize in Chemistry was divided equally between Walter Kohn "for his development of the density-functional theory" and John A. Pople "for his development of computational methods in quantum chemistry".

Taking a legacy code to high-performance computing levels is not a simple task. Rather than introducing optimizations at the lowest level (i.e. single lines of code), it is better to step back and identify *kernels*. That is, functions, loops or in general code structures which take up a sizable amount of computational workload. The resulting modular code allows for simpler optimizations, code maintainability, and extensibility. If the code itself does not expose isolated kernels, identifying them is the first step that needs to be performed in order to yield performance improvements of lasting value. To this end, we follow an approach which begins with the mathematical formulation of the physical problem, singles out efficient data structures, and maps the mathematical operations onto portable high-performance kernels. While our strategy can resemble previous similar attempts [5, 6], the end result of this work is a road map granting the revived code access to multi- and many-core architectures and enabling the simulation of large-scale materials.

The paper is organized in four main sections. Sec. 2 introduces the reader to basic concepts of quantum mechanics, the fundamentals of Density Functional Theory (DFT) and the mathematical setup leading to the initialization of the Hamiltonian and Overlap matrices. This section, with the exception of Sec. 2.3, can be skipped by the reader familiar with DFT and its many flavors. The following section offers a brief overview of the FLEUR code and some insights on the algorithmic strategies implemented there. In Sec. 4 we present the core of our original contribution including the rationale behind the specific choices of algorithms and memory layout. This section is quite technical and, at the same time, dense with details we hope to be quite useful to the dedicated developer. The last section is devoted to numerical results and performance measurements.

2. The FLAPW method

The FLEUR code is based on the widely accepted Density Functional Theory [7, 8] theoretical framework. This theory has been and is currently used to simulate physical properties of materials used in the development of devices such as Blu-ray discs, memory chips, photovoltaic cells, just to name a few. There exists a wide variety of approaches that can be used to "translate" the DFT mathematical layout into a computational tool. In this work, we focus on the Full-potential Linearized Augmented Plane Wave (FLAPW) variant [9, 10], one of the most accurate methods due to its particular discretization of the DFT fundamental equations. In contrast to others variants that use only an effective potential describing the dynamics of the valence electrons, FLAPW is an all-electron method. This means that on the flip side, FLAPW explicitly describes all of the (potentially large number of) electrons in the material with a much larger number of basis functions and consequently is a quite computationally expensive method.

In this section, we focus on the mathematical structure of DFT, and give a cursory overview of its physical foundation. The material presented is meant for the reader unfamiliar with quantum mechanics and provides a rough overview of the mathematical model and terminology behind DFT in general and the FLAPW method in particular. For the sake of clarity we start with the concept of *wave functions* and introduce the Schrödinger equation before proceeding to a very short overview of what constitutes the DFT formalism.

At the theoretical level, the quantum mechanical description of an atomic or molecular physical system is given by a complex-valued wave function Ψ, which expresses the probability amplitude to find an electron in a region of space $\Omega \subset \mathbb{R}^3$. In the case of stationary multi-atomic

systems, Ψ is a high-dimensional function

$$\Psi\left(\mathbf{r_1}, \mathbf{s_1}; \ldots; \mathbf{r_{N_e}}, \mathbf{s_1}\right) : \left(\mathbb{R}^3 \times \left\{\pm\frac{1}{2}\right\}\right)^{N_e} \longrightarrow \mathbb{C}.$$

Ψ describes the electron's dynamics[2] and is the solution of the time-independent Schrödinger equation

$$\hat{H}\Psi \coloneqq \left[-\sum_{i=i}^{N_e} \frac{\hbar^2}{2m_e}\nabla^2_{\mathbf{r_i}} + V\left(\{\mathbf{r_i}\}\right)\right] \Psi\left(\{\mathbf{r_i}, \mathbf{s_i}\}\right) = E\Psi\left(\{\mathbf{r_i}, \mathbf{s_i}\}\right), \tag{1}$$

where N_e is the total number of electrons whose positions and spins are characterized by the set of variables $\{\mathbf{r_i}, \mathbf{s_i}\} \equiv \mathbf{r_1}, \mathbf{s_1}; \ldots; \mathbf{r_{N_e}}, \mathbf{s_{N_e}}$. The operator $\hat{H}$ on the LHS of Eq. (1) is the Hamiltonian of the physical system, and represents its total energy E appearing on the equation's RHS. As such, Eq. (1) is an eigenvalue equation which, once discretized, becomes an algebraic eigenproblem (see Sec. 2.3).

The solution of the equation usually depends on a large set of both discrete-valued and continuous-valued parameters, and encodes the probabilistic behavior of all the degrees of freedom involved (positions, spins, momenta, etc.). Already for systems with more than two electrons, the exact solution of such an equation is quite challenging. For more than two electrons, solving the Schrödinger equation is known to hit the so called Exponential Wall problem[3]: not only does the time to solution increase exponentially, but storing a solution requires more memory than the total number of subatomic particles in the universe [**?**].

In the case of complex multi-atomic systems, only approximate solutions exist. One of the most successful frameworks for approximate solutions is Density Functional Theory. DFT constrains the types of available solutions to the ground states of quantum systems. Despite its apparent limitation, DFT proved to be extremely successful, is the subject of thousands of scientific papers every year [11], and is widely used in both Quantum Chemistry and Materials Science computations.

2.1. Density Functional Theory in a nutshell

DFT is based on the fundamental work of Hohenberg and Kohn [12], and successive extension by Kohn and Sham [13]. By establishing a one-to-one correspondence between the electronic charge density $n\left(\mathbf{r}\right)$ and the total potential V, the Hohenberg-Kohn theorem moves away from a quantum mechanical description using the wavefunction Ψ,[4] to the more manageable one-particle charge density $n(\mathbf{r}) : \mathbb{R}^{\mathbf{3}} \to \mathbb{R}$

$$n\left(\mathbf{r}\right) = N_e \cdot \int \ldots \int \Psi\left(\mathbf{r}, \mathbf{r_2}, \ldots, \mathbf{r_{N_e}}\right)^* \Psi\left(\mathbf{r}, \mathbf{r_2}, \ldots, \mathbf{r_{N_e}}\right) \mathrm{d}\mathbf{r_2} \ldots \mathrm{d}\mathbf{r_{N_e}}. \tag{2}$$

Such a shift in the description of the quantum system implies a reduction of degrees of freedom from the $3N$ to just 3, which is one of the main reasons DFT is so appealing.

[2] We are implicitly assuming to be in the realm of the validity of the Born-Hoppenheimer adiabatic approximation.

[3] Also known as Van Vleck catastrophe.

[4] In the following treatment of DFT, we omit any mention to the spin variable to avoid delving into the full relativistic treatment of the equations.

Building on the Hohenberg and Kohn theorem, Kohn and Sham showed that it is possible to reformulate the initial high-dimensional Schrödinger equation in terms of $N > N_e$ one-dimensional Schrödinger-like equations

$$\hat{H}_{\mathrm{KS}}\,\psi_i\left(\mathbf{r}\right) = \left[-\frac{\hbar^2}{2m_e}\nabla^2_{\mathbf{r}} + V_{\mathrm{eff}}[n]\left(\mathbf{r}\right)\right]\psi_i\left(\mathbf{r}\right) = \epsilon_i\psi_i\left(\mathbf{r}\right). \tag{3}$$

The initial potential $V(\{\mathbf{r}\})$ is substituted by an effective potential $V_{\mathrm{eff}}\left[n\right]$

$$V_{\mathrm{eff}}\left[n\right] = V_{\mathrm{ext}}\left(\mathbf{r}\right) + \int \frac{n\left(\mathbf{r}'\right)}{\left|\mathbf{r}-\mathbf{r}'\right|}\mathrm{d}\mathbf{r}' + V_{\mathrm{xc}}\left[n\right]\left(\mathbf{r}\right) \tag{4}$$

that depends functionally on the charge density $n\left(\mathbf{r}\right)$ with the sole exception of V_{ext}, which is the nuclei's Coulomb term. The charge density is now computed as a function of all ψ_i

$$n\left(\mathbf{r}\right) = \sum_{i=1}^{N_e}\left|\psi_i\left(\mathbf{r}\right)\right|^2, \tag{5}$$

where the sum is intended over the lowest N_e eigenvalues ϵ_i. Notice that the functions ψ_i and the eigenvalues ϵ_i in Eq. (3) do not have any direct physical interpretation. The ψ_is just determine the physical density $n\left(\mathbf{r}\right)$ of Eq. (5) and the ϵ_is contribute, together with other terms, to the total energy E. While apparently simple, the intricacies of the theory are hidden in the explicit expression of the exchange-correlation potential V_{xc} [8].

Looking at the complete set of Kohn-Sham equations (3)–(5), one realizes there is a cyclic dependence. The $\psi_i\left(\mathbf{r}\right) : \mathbb{R}^3 \to \mathbb{C}$ are the solutions to Eq. (3), which cannot be solved without first calculating V_{eff} of Eq. (4). On the other hand, V_{eff} is dependent on the electron density $n\left(\mathbf{r}\right)$, which requires a set of valid ψ_i to be calculated in the first place. For this reason, Equations (3) are said to be non-linearly coupled.

The usual procedure to resolve the dilemma is a *self-consistency* approach: one starts from an electron density $n\left(\mathbf{r}\right)_{\mathrm{start}}$ derived from inexpensive yet somewhat accurate wave function calculations,[5] computes an effective potential $V_{\mathrm{eff}}\left[n\right]$, and solves Eq. (3). The resulting ψ_i's and eigenvalues ϵ_i's are then used to compute a new density as in Eq. (5), which is compared to the initial one. If the two densities disagree, the self-consistent cycle is repeated with an opportunely modified charge density. Once the density difference converges below some defined threshold, the procedure is stopped.

So far we have not presented a particular method for solving Eq. (3). This is where the various "flavors" of DFT differ. The concept of FLAPW distinguishes itself by the particular discretization that is chosen for the Kohn-Sham equations. From physical observations, the wave functions are known to have different symmetries in distinct regions of the space: close to the atomic nuclei, solutions tend to be spherically symmetric and strongly varying, while further away from the nuclei, they can be approximated as almost constant and lack this symmetry. Defining a cutoff distance for these structurally different solution regions leads to a landscape composed of non-overlapping spheres (called *muffin tins*, MT) separated by *interstitial* (INT) areas. Refining this concept and developing a rigorous model for describing

[5] The initial density needs to be in the convex hull of the converged density.

ψ_i and the potential $V_{\text{eff}}[n]$ in these distinct areas defines the FLAPW method [9, 10, 14].

2.2. The Full-potential Linearized Augmented Plane Wave method

The distinct flavors of the Kohn-Sham approach stem from the challenge of finding a representation that allows for a convenient numerical construction of the single-particle solutions ψ_i. FLAPW is a basis function method, which means that it expands ψ_i with functions φ_t exhibiting enough adjustable parameters to accurately represent the physical system with as few N_G basis functions as possible:

$$\psi_i(\mathbf{r}) = \sum_{t=1}^{N_G} c_{t,i}\, \varphi_t(\mathbf{r}). \tag{6}$$

This concept is not unlike the one encountered in a Fourier transformation, where a function $f(x)$ is represented by a sum over complex functions $e^{i\omega x}$, thereby shifting the problem from finding the unknown functions ψ_i to determining the unknown coefficients $c_{t,i}$.

As already mentioned, the FLAPW idea is to split the basis function φ_t into two parts. In the MT regions, φ_t is determined by considering only the spherically symmetric part of the potential. Each region of space is then treated as a standalone system with its own, simpler Schrödinger equation, which is both one-dimensional and spherically symmetric. Its solution, $u_{l,a}(r)$, then serves as a localized basis function for each distinct atom labeled by the index a:

$$\left[-\frac{\hbar^2}{2m}\frac{\partial^2}{\partial r^2} + \frac{\hbar^2}{2m}\frac{l(l+1)}{r^2} + V_{\text{eff(ph)}}(r) - E_l\right] r u_{l,a}(r) = 0. \tag{7}$$

In the MTs, the full functions φ_t are then given by a combination of *radial* functions $u_{l,a}(r)$ and *spherical harmonics* $Y_{l,m}(\hat{\mathbf{r}})$; the former only depend on the distance from the MT center, while the latter are just functions of the MT spherical angles. The subscripts l, m denote independent solutions.

In the interstitial region, where the potential varies more slowly than in the MT regions, plane waves constitute an excellent basis set. Plane waves are written as complex exponentials $\exp(i\mathbf{k}\cdot\mathbf{r})$, where $\mathbf{k}$ is the *wave vector* that points in the direction the wave propagates and plays exactly the same role as ω in a Fourier transform. In DFT, the physical systems of interest are, for the most part, crystals and can be represented by a *lattice*, i.e. a periodic arrangement of atoms on a discrete spatial grid. This periodicity in real space is reflected in a similar periodicity in *momentum* space, where the Bloch theorem prescribes $\mathbf{k}$ to have independent values only in the Brillouin zone, a unit cell in momentum space [15].

In the end, the resulting ansatz for the basis functions φ_t brings together spherical contributions from the muffin tin spheres and plane waves from the interstitial zone. Overall, the basis size N_G determines the number of expansion coefficients $c_{t,i}$ to be stored per wave function ψ_i, and as such constitutes a truncated expansion which is usually referred to as *discretization*. The complete basis functions φ_t are given by a piece-wise definition on each of the N_{A} MT[6]

[6] One for each atom.

and the surrounding INT regions:

$$\varphi_t(\mathbf{r}) = \begin{cases} \sum_{l=0}^{l=l_{\max}} \sum_{m=-l}^{m=+l} \left[A_{(l,m),a,t} u_{l,a}(r) + B_{(l,m),a,t} \dot{u}_{l,a}(r)\right] Y_{l,m}(\hat{\mathbf{r}}_a) & a^{\text{th}}\ MT \\ \frac{1}{\sqrt{\Omega}} \exp(i\mathbf{K}_t \cdot \mathbf{r}) & INT \end{cases} \tag{8}$$

Both the coefficients $A, B \in \mathbb{C}$ are necessary to guarantee $\varphi_t \in C^1$. A and B have three independent dimensions: For each index t and each atomic index a, the function φ_t needs to satisfy the requirement of being of class C^1 for all values of $L \equiv (l, m)$, thus a different coefficient is needed for every index tuple (L, a, t).[7]

The functions $u_{l,a}(r)$ appear together with their energy derivatives $\dot{u}_{l,a} := \frac{\partial u_{l,a}}{\partial E_l}$ so that, even when E_l is kept fixed, there is enough variational freedom to obtain accurate results for the band energies. The spherical harmonics $Y_{lm}(\hat{\mathbf{r}})$ form a complete basis on the unit sphere ($\hat{\mathbf{r}} = \mathbf{r}/|\mathbf{r}|$), and capture the angular part of the exact solution for a particle in the field of a single independent atom. We abbreviate $L = (l, m)$, so it is important to distinguish between a capital L referring to both l and m, and a lower-case l, which is just the first component of the tuple. t, as used above, ranges over the size of the nonequivalent plane wave functions in INT, and is used to label the vector $\mathbf{G}_t$ living in the space reciprocal to $\mathbf{r}$. The sum of the vector $\mathbf{G}_t$ with the momentum $\mathbf{k}$, which lives in the Brillouin zone, characterizes the specific wave function entering in the basis set. To avoid cluttering, such a sum is condensed in the definition of the vector $\mathbf{K}_t$. The finite size of the basis set is determined by imposing a cutoff value $\mathbf{K}_{\max} \geq \mathbf{K}_t = \mathbf{k} + \mathbf{G}_t$ on $\mathbf{K}_t$. In other words, by choosing a cutoff value $\mathbf{K}_{\max}$, we indirectly affect the total number of $\mathbf{G}_t$ available per choice of $\mathbf{k}$-point, which dictates the size of the basis set N_G. Ω is the volume of the simulation cell and is introduced so as to satisfy the normalization to 1 of the wave functions' square (i.e. the probability density).

2.3. Constructing the Hamiltonian and overlap matrices

Having introduced the reader to the details of the FLAPW method, we are finally in a position to define the operators whose numerical initialization is at the heart of this paper. We start by plugging Eq. (6) into Eq. (3) followed by a left-multiplication by a complex conjugate basis function $\varphi^*_{t'}(\mathbf{r})$ labeled by the index t'. Integrating the resulting equation over the distinct regions of space transforms the Kohn-Sham equations into an algebraic *generalized eigenproblem* for the coefficients $\mathbf{c}_i = (c_{1,i} \dots c_{N_G,i})^T$,

$$\sum_{t=1}^{N_G} (H)_{t',t}\, c_{t,i} = \epsilon_i \sum_{t=1}^{N_G} (S)_{t',t}\, c_{t,i} \qquad \Longrightarrow \qquad \mathbf{H} \cdot \mathbf{c}_i = \epsilon_i\, \mathbf{S} \cdot \mathbf{c}_i \tag{9}$$

with the entries of the Hamiltonian and Overlap matrices — respectively $\mathbf{H}$ and $\mathbf{S}$ — given by

$$(H)_{t',t} = \sum_a \iint \varphi^*_{t'}(\mathbf{r}) \hat{H}_{\text{KS}} \varphi_t(\mathbf{r}) \mathrm{d}\mathbf{r}, \quad (S)_{t',t} = \sum_a \iint \varphi^*_{t'}(\mathbf{r}) \varphi_t(\mathbf{r}) \mathrm{d}\mathbf{r}. \tag{10}$$

In order to have an explicit formulation of the $\mathbf{H}$ and $\mathbf{S}$ matrices, we substitute Eq. (8) in

[7] See Appendix A.1 for more details.

Eq. (10) and focus solely on the MT regions. It is these regions where the initialization of the Hamiltonian and Overlap matrices is by far the most computationally intensive task. On the contrary, the interstitial part is simple and some of the contributions can even be computed analytically. The step from Eq. (9) to a workable expression exploits the properties of the basis functions and yields final expressions directly depending on the set of A and B coefficients.

The Overlap matrix

$$(S)_{t',t} = \sum_{a} \sum_{L=(l,m)} A^*_{L,a,t'} A_{L,a,t} + B^*_{L,a,t'} B_{L,a,t} \left\| \dot{u}_{l,a} \right\|^2 \tag{11}$$

is obtained by exploiting the mutual orthogonality of the spherical harmonics $Y_{l,m}(\hat{\mathbf{r}}_a)$ and by enforcing the von Neumann condition on the radial functions ($\int u^*_{l,a}(r)\, \dot{u}_{l,a}(r) = 0$) on each MT region separately. Due to the non-orthornormality of the basis function set (8), the matrix $\mathbf{S}$ is non-diagonal and effectively dense. On the up side, due to the positivity requirement of the probability density $\iint \varphi^*_t \varphi_{t'} > 0$, the Overlap matrix is Hermitian and positive definite. On the down side, the basis function set is by definition overcomplete, and some of the functions in the set may almost depend on the others. Such dependency may be the source of possible rank deficiencies leading to an ill-conditioned $\mathbf{S}$ matrix having few singular values close to zero.

The Hamiltonian matrix is given by:

$$\begin{aligned}(H)_{t',t} = \sum_{a} \sum_{L',L} &\left(A^*_{L',a,t'}\, T^{[AA]}_{L',L;a}\, A_{L,a,t} \right) + \left(A^*_{L',a,t'}\, T^{[AB]}_{L',L;a}\, B_{L,a,t} \right) \\ &+ \left(B^*_{L',a,t'}\, T^{[BA]}_{L',L;a}\, A_{L,a,t} \right) + \left(B^*_{L',a,t'}\, T^{[BB]}_{L',L;a}\, B_{L,a,t} \right).\end{aligned} \tag{12}$$

The new matrices $T^{[\dots]}_{L',L;a} \in \mathbb{C}^{N_L \times N_L}$ are dense as well and their computation involves multiple integrals between the basis functions and the non-spherical part of the potential V_{eff}.[8] Their size depends on the cutoff over the spherical angular momentum which, in turn, is contingent on the specific atom they are associated with. Overall the Hamiltonian is also Hermitian but is indefinite and presents always some negative eigenvalues corresponding to bounded states.

It needs to be noted that the set of basis functions in Eq. (8) are implicitly labeled by the values the variable $\mathbf{k}$ takes in the Brillouin zone. Not only is this dependence embedded in the definition of $\mathbf{K}_t$, but it also appears in the definition of the coefficients A and B (see Appendix A.1). Consequently there are multiple Hamiltonian and Overlap matrices, one for each independent $\mathbf{k}$-point.

We end this section with a brief digression on the computational cost of one full self-consistent cycle. In the FLAPW method, this cycle can be broken up into the following steps:

1. An initial charge density $n(\mathbf{r})_{\text{start}}$ is used to compute the potential $V_{\text{eff}}[n(\mathbf{r})]$ [Eq. (4)];
2. The spherical part of V_{eff} is used to compute the radial functions $u_{l,a}$ [Eq. (7)] which are then used to derive the A, B coefficients [Eqs. (A.1) and (A.2)];
3. Hamiltonian $\mathbf{H}$ and Overlap $\mathbf{S}$ matrices are initialized [Eqs. (12) and (11)];
4. The generalized eigenvalue problems $\mathbf{H} \cdot \mathbf{c}_i = \epsilon_i\, \mathbf{S} \cdot \mathbf{c}_i$ are solved numerically to return values ϵ_i and vectors of coefficients $\mathbf{c}_i$, which are then used to calculate a new charge density $n(\mathbf{r})$ [Eq. (5)];

[8] See Appendix A.2.

5. If self-consistency is not reached, a charge density mixing scheme is invoked before starting a new cycle.

Out of all the steps above, steps 3 and 4 account for more than 80% of CPU time. Having cubic complexity $\mathcal{O}\left((N_G)^3\right)$, the eigenproblem solution is usually considered the most expensive of the two. It turns out that generating the matrices may be as expensive. Let us define with N_A and N_L the range of the summations $\sum_a$ and $\sum_L$ respectively. Then, a back-of-the-envelope estimate shows that Eqs. (11) and (12) have complexity equal to $\mathcal{O}\left(N_A \cdot N_L \cdot (N_G)^2\right)$ and $\mathcal{O}\left(N_A \cdot N_L \cdot N_G \cdot (N_L + N_G)\right)$ respectively. A typical simulation uses approximately N_G basis functions, with N_G ranging from about $50 \cdot N_A$ to about $80 \cdot N_A$, and an angular momentum $l_{\max} \leq 10$, which results in $N_L = (l_{\max} + 1)^2 \leq 121$. It follows that the factor $N_A \cdot N_L$ is roughly of the same order of magnitude as N_G so that the generation of $\mathbf{H}$ and $\mathbf{S}$ also displays cubic complexity $\mathcal{O}\left((N_G)^3\right)$. In the reminder of this paper, we focus on the implementation of Eqs. (11) and (12) within the FLEUR software, and illustrate how the traditional implementation can be re-engineered and optimized to take advantage of Basic Linear Algebra Subroutines (BLAS). For the reader interested in improving the computational aspects of the eigenproblem solution in FLAPW, we refer to [19, 20, 21].

3. The FLEUR code

The FLEUR code family is a software project [1] for the computation of ground state and excited state properties of solids. FLEUR supports calculations on a plethora of different system types, and is particularly renowned for the simulation of non-collinear magnetic systems as well as thin-film geometries. The entire package was developed over the course of 20+ years at the Peter Grünberg Institute within the Forschungszentrum Jülich. It is a full blown DFT code based on the FLAPW method with more than 100,000 lines of code distributed on more than 500 routines. Initially written in Fortran 77 and later partially modernized by introducing concepts of Fortran 90, FLEUR was not designed with high-performance computing as the number one priority in mind. Eventually, such a decision has lead to a software design with undesirable properties which makes it hard to adapt FLEUR to modern parallel architectures.

One of the most relevant strategic choices in the multi-years implementation of FLEUR was the minimization of its memory footprint. This choice is easily understood by looking back at the computing architectures available at the turning of the last century and comparing them with the memory impact of an average-size simulation of FLEUR. A simple rough estimate based on the typical size of the involved mathematical objects shows why this is the case. We have seen at the end of Sec. 2 that we roughly need to use N_G basis functions, with N_G between $50 \cdot N_A$ and $80 \cdot N_A$, while $N_L \leq 121$. Even for a fairly small system with ~ 100 atoms, each of the matrices $(H)_{t',t}$ and $(S)_{t',t}$ would have size $N_G^2 \geq (50 \cdot 100)^2 \,\widehat{=}\, 0.38\,\text{GiB}$. Similarly, the A, B tensors of this system would each have a size of $N_L \cdot N_A \cdot N_G \sim 100 \cdot 100 \cdot (50 \cdot 100) \,\widehat{=}\, 0.76\,\text{GiB}$. Storing these objects in memory explicitly for each $\mathbf{k}$-point would have soon outgrown the memory per node available on the Jülich cluster.[9] By choosing to minimize the memory footprint, the FLEUR developers avoided to run into the typical memory contraints of these

[9] For instance the CRAY SV1ex, which was in operation at the Forschungszentrum Jülich between 1996 and 2002, had 2 GiB of memory per CPU. IBM Blue Gene/Q, which is the current leading platform, has only 1 GiB per core.

early computing platforms at the cost of introducing rigid data structures. In order to reap the real benefit from parallelization, such rigid software structures have to undergo drastic changes.

In scientific computing it is commonly accepted practice to layer modules where the main computational operations are provided by the lowest kernels. Such kernels have usually clearly defined input and output quantities while the higher layers are unconcerned with the specific tasks carried on by the lowest kernels. For instance, this is the philosophy behind the use of the BLAS library which allows for a flexible implementation and a simple structure that makes the correctness of execution easy to verify.

The layout of the FLEUR code does not follow such accepted practice. Specifically, the construction of the full **H** and **S** matrices lacks the modern coding practice of encapsulating different functionalities in a set of well defined layers. For example, the module in which the spherical part of the Hamiltonian is computed is the same module where also the Overlap matrix is initialized. Contributions coming from the non-spherical part of the potential are computed in a separate module and added on top of the spherical part. When FLEUR is used to simulate magnetic systems, electronic spins and localized orbitals are incorporated by calling a number of other modules. While seemingly modular, the FLEUR main computational "kernel" puts together the results of all these computations in a routine spanning $\sim$ 1,500 lines of practically undocumented code with many dozens of cryptically named global variables. The resulting code is quite challenging to understand and optimize to say the least.

3.1. FLEUR's algorithm for the **H** *and* **S** *matrices*

In the rest of this section, we briefly outline the core algorithmic choices used in FLEUR to implement the generation of **H** and **S**. When implementing the computation of Eq. (12), the FLEUR code almost never uses external libraries and implements matrix multiplications using explicit loops without blocking. In practice all computations are performed in entry-wise fashion and each summation is "translated" in as many nested loops.

Another peculiarity of the FLEUR code which is worth mentioning is due to the different contributions to the T-matrices. The diagonal terms — which FLEUR internally does not consider as part of the T-matrices, but generates in a separate loop — are needed up to a specific cutoff value l_{sph}. This choice implies that there are $L_{\text{sph}} = (l_{\text{sph}} + 1)^2 \equiv N_{\text{sph}}$ total entries to consider. On the other hand, the non-spherical contributions are only needed up to a smaller cutoff value l_{nonsph}, so this part of the matrix has dimension $L_{\text{nonsph}} = (l_{\text{nonsph}} + 1)^2$. Only this latest contribution to the T-matrices is dense. Storing both contributions in a single matrix results in the structure shown in Fig. 1. For a realistic choice of parameters, i.e. $l_{\text{sph}} = 8$ and $l_{\text{nonsph}} = 6$, about half of the entries can be zero.

The FLEUR code exploits the structure of the T-matrices by storing separately the spherical and the non-spherical part of the matrices. While this choice minimizes FLEUR's memory footprint, it is realized by disregarding the matrix structure of Eq. (12) and effectively leading to a non-efficient implementation. In our reengineering of the FLAPW algorithm we pay the price of a bigger memory footprint but maintain the full structure of the T-matrices so as to exploit the full potential of level 3 BLAS routines. One could argue that extra performance could be achieved by considering a more fine-grained structure where each multiplication with $T_{L,L'}$ will be split in a full matrix-matrix product for the dense part and a matrix-vector product for the diagonal part. We leave this further optimization to future work.

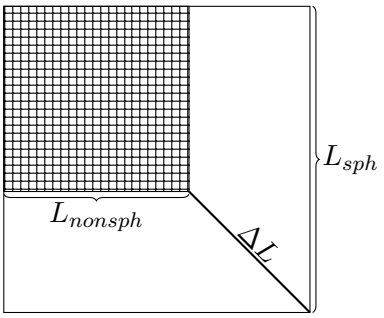

Figure 1: Structure of a T-matrix. Outside of the top-left dense area, all entries except the diagonal are zero. The size of the lower right submatrix that is diagonal is $\Delta L = L_{\text{sph}} - L_{\text{nonsph}}$.

In the implementation of Eq. (11), FLEUR developers used the same philosophy followed in the implementation of Eq. (12): Each sum is realized as a number of nested loops with little or no use of kernels from specialized libraries. Despite the lack of a high-performance approach, some clever mathematical manipulations based on the properties of the spherical harmonics $Y_{l,m}$ are used to reduce the overall amount of computation needed. This is a good example of the ingenuity of the FLEUR developers.

For the sake of simplicity, we restrict the analysis to the part of Eq. (11) dealing with A coefficients

$$M_{t',t} = \sum_{a} \sum_{L=(l,m)} A^*_{L,a,t'} A_{L,a,t}. \tag{13}$$

Matching the INT and MT part of φ_t at their boundary results in coefficients A expressed in abbreviated form as the multiplication of a prefactor with a (real-valued) term $f_{l,a,t}$ consisting of the matching radial and Bessel functions evaluated at the boundary[10]

$$A_{L,a,t} = \frac{4\pi}{W_{l,a}\sqrt{\Omega}} i^l \exp\left(i\mathbf{K}_t \cdot \mathbf{x}_a\right) Y^*_{l,m}\left(\mathbf{R}_a \hat{K}_t\right) \cdot f_{l,a,t}.$$

Now Eq. (13) becomes

$$M_{t',t} = \sum_{L,a} \left(\frac{4\pi}{W_{l,a}\sqrt{\Omega}} i^l \exp\left(i\mathbf{K}_{t'} \cdot \mathbf{x}_a\right) Y^*_{l,m}\left(\mathbf{R}_a \hat{K}_{t'}\right) \right)^* \frac{4\pi}{W_{l,a}\sqrt{\Omega}} i^l \exp\left(i\mathbf{K}_t \cdot \mathbf{x}_a\right) Y^*_{l,m}\left(\mathbf{R}_a \hat{K}_t\right) \cdot f_{l,a,t'} f_{l,a,t}.$$

We now collect similar terms, simplify the imaginary units $\left(i^l\right)^* i^l = 1$ and write out the

[10] See Eq. (A.1) for details.

sum over $L = (l, m)$ explicitly

$$M_{t',t} = \frac{(4\pi)^2}{\Omega} \sum_a \exp\left[i\left(\mathbf{K}_t - \mathbf{K}_{t'}\right)\cdot \mathbf{x}_a\right] \sum_{l=0}^{l_{\text{sph}}} \frac{f_{l,a,t'} f_{l,a,t}}{W_{l,a}^2} \underbrace{\sum_{m=-l}^{l} Y_{l,m}^*\left(\mathbf{R}_a \hat{K}_t\right) Y_{l,m}\left(\mathbf{R}_a \hat{K}_{t'}\right)}_{\star}.$$

The sum over m can be removed by using the well known identity $P_l\left(\hat{K}_t \cdot \hat{K}_{t'}\right) \frac{2l+1}{4\pi} = \star$, relating the Legendre polynomial P_l with the spherical harmonics $Y_{l,m}$. The resulting expression

$$M_{t',t} = \frac{(4\pi)^2}{\Omega} \sum_a \exp\left[i\left(\mathbf{K}_t - \mathbf{K}_{t'}\right)\cdot \mathbf{x}_a\right] \sum_{l=0}^{l_{\text{sph}}} \frac{f_{l,a,t'} f_{l,a,t}}{W_{l,a}^2} \cdot (2l+1)\, P_l\left(\hat{K}_t \cdot \hat{K}_{t'}\right) \tag{14}$$

maintains the same matrix-matrix product structure, but with a reduced size: the object $A_{L,a,t}$ is replaced by $f_{l,a,t}$, which has a smaller first dimension. Initially, for a given l, there were $2l+1$ values of m, resulting in a total of $N_L = (l_{\text{sph}}+1)^2$ terms in the sum in Eq. (13). By removing the sum over m there remain only $l_{\text{sph}}+1$ elements, yielding a reduction by a factor of $(l_{\text{sph}}+1) \sim 10$ for the typical values of $l_{\text{sph}} \sim 8, \ldots, 10$. The same simplification carries over to the part of $\mathbf{S}$ involving the B coefficients, since they share the same prefactor with the A coefficients. FLEUR also structures computation differently by precalculating the phase factors for each atom type and thereby replacing the sum over atoms by a sum over types, though it does not specifically recognize the matrix-matrix product as such and does not use BLAS calls.

Out of this digression we can take home an important message. The clever manipulation in the initialization of the matrix $\mathbf{S}$ operated in FLEUR is hidden in the implementation of the routine and only partially documented. As such, it is a typical example of "premature optimization" [22] which is quite hard to spot and renders the re-working and optimization of the code by an expert programmer a very hard task. On the bright side, this type of reduction in complexity — which preserves the matrix structure of the operands — once identified can be included in further algorithmic optimizations beyond the ones described in the next section of this work.

4. Optimized generation of H and S through Dense Linear Algebra: HSDLA

In this section, we present HSDLA, our algorithm for the computation of the Overlap matrix $\mathbf{S}$ (see Eq. (11)) and the Hamiltonian $\mathbf{H}$ (see Eq. (12)), which accounts for roughly 40% of FLEUR's execution time. Our starting point for this computation are the coefficients $A_{L,a,t}$ and $B_{L,a,t}$, which are computed by an efficient implementation of Eq. (A.1) and (A.2), and the $T^{[\ldots]}_{L,L',a}$ and $\dot{u}_{l,a}$, which are extracted directly from FLEUR. Since from this point on all calculations are essentially linear algebra operations, we will treat all involved objects as matrices, discarding bold fonts and dropping the indices L, and t: $A_a, B_a \in \mathbb{C}^{N_L \times N_G}$, $T_a^{[\ldots]} \in \mathbb{C}^{N_L \times N_L}$, and $H, S \in \mathbb{C}^{N_G \times N_G}$. In terms of these matrices and the diagonal matrix

$\dot{U}_a \in \mathbb{C}^{N_L \times N_L}$ with entries $(\dot{U}_a)_{(l,m),(l,m)} = \dot{u}_{l,a}$, Eqs. (12) and (11) become:

$$S \coloneqq \sum_{a=1}^{N_A} A_a^H A_a + B_a^H \dot{U}_a^H \dot{U}_a B_a, \tag{15}$$

$$H \coloneqq \sum_{a=1}^{N_A} A_a^H T_a^{[AA]} A_a + A_a^H T_a^{[AB]} B_a + B_a^H T_a^{[BA]} A_a + B_a^H T_a^{[BB]} B_a. \tag{16}$$

In subsection 4.1, we give an overview of the employed libraries, their implied storage formats, and the principles guiding our optimization. In order to understand its rationale, we present below how we plan to subdivide the computation of S and H:

$$S \coloneqq \underbrace{\sum_{a=1}^{N_A} A_a^H A_a}_{S_{AA}} + \underbrace{\sum_{a=1}^{N_A} B_a^H \dot{U}_a^H \dot{U}_a B_a}_{S_{BB}},$$

$$H \coloneqq \underbrace{\sum_{a=1}^{N_A} A_a^H T_a^{[AA]} A_a}_{H_{AA}} + \underbrace{\sum_{a=1}^{N_A} A_a^H T_a^{[AB]} B_a + B_a^H T_a^{[BA]} A_a}_{H_{AB+BA}} + \underbrace{\sum_{a=1}^{N_A} B_a^H T_a^{[BB]} B_a}_{H_{BB}}.$$

Starting with $S, H \coloneqq 0 \in \mathbb{C}^{N_G \times N_G}$, these five contributions are treated separately as follows:

- subsection 4.2 constructs $S \mathrel{+}= S_{AA}$ and introduces the used memory layout,
- subsection 4.3 extends this to $S = S_{AA} + S_{BB}$,
- subsection 4.4 computes $H \mathrel{+}= H_{AB+BA}$,
- subsection 4.5 incorporates H_{BB} to obtain $H = H_{AB+BA} + H_{BB}$, and
- subsection 4.6 is concerned with the update $H \mathrel{+}= H_{AA}$.

While up to subsection 4.6 we work with the simplifying assumption that all $T_a^{[\ldots]}$ are of size $N_L \times N_L$, subsection 4.7 describes the changes needed in the developed algorithms to account for variations in the size of these matrices. After this point, we have isolated implementations of the updates $S \mathrel{+}= S_{AA} + S_{BB}$, $H \mathrel{+}= H_{AB+BA} + H_{BB}$, and $H \mathrel{+}= H_{AA}$; aiming at minimizing HSDLA's memory footprint, in subsection 4.8, we combine these components into our final algorithm.

4.1. Using high performance BLAS and LAPACK

It is our goal to compute S and H with the high performance BLAS (Basic Linear Algebra Subprograms) and LAPACK (Linear Algebra PACKage) libraries. Using these libraries' standardized APIs, we can directly benefit from the performance of highly optimized implementations (e.g., Intel's Math Kernel Library), which commonly reach 80% – 90% of a computer's efficiency in terms of available FLOPs/s (floating point operations per second) both sequentially and across multiple threads.

While BLAS and LAPACK provide a wide range of basic building blocks for dense linear algebra operations, complex computations such as those of H and S do not map directly

to any of them and have to be further decomposed. To make the most of BLAS and LAPACK's high performance, we apply the following optimization guidelines (in decreasing order of importance).

- Cast as much computation in terms of BLAS and LAPACK routines. On a single core, these libraries are at least 10 times faster than naïve implementations; on multi-core systems with shared memory the speedup is even larger.
- Reduce the amount of computation. This is for instance achieved by avoiding redundant operations, combining operations mathematically, and only computing the lower triangular portion of Hermitian matrices.
- Combine small operations into few large operations. Optimized linear algebra libraries generally reach higher performance for larger operations, especially when using many threads.
- Reduce the memory footprint. Since the main memory of modern computers is considerably larger than in the early days of FLEUR, this goal is secondary to the efficiency-related targets above.

Storage format. Both BLAS and LAPACK work with the same data layout, where complex matrices are stored element-wise "by column" (column-major order).

- Each complex number is stored as two consecutive double precision floating point numbers, respectively representing its real and complex component.
- The numbers in a column of the matrix are stored consecutively in memory.
- The columns of a matrix are stored with a constant stride, known as the matrix's leading dimension. This stride, which is the offset in memory between two elements in the same matrix row can be the height of the matrix or larger.

Our input matrices A_a, B_a and the $T_a^{[\dots]}$ are generated in this format, where, for now, the leading dimensions are simply the height of the matrices.

4.2. Computing S: Constructing S_{AA}

We begin the computation of S with

$$S \coloneqq S_{AA} = \sum_{a=1}^{N_A} A_a^H A_a. \tag{17}$$

First off, both S and the contributions S_{AA} are Hermitian. We make use of this property by only computing the lower triangular half of S_{AA} (including the diagonal). The BLAS library provides the kernel `zherk` that performs the required updates $S \mathrel{+}= A_a^H A_a$. A basic algorithm that uses this kernel to compute Eq. (17) follows.

```
1  for a := 1, ..., N_A:
2    S += A_a^H A_a                          (zherk: 4 N_L N_G^2 FLOPs)
```

Listing 1: $S \mathrel{+}= S_{AA}$ by N_A `zherk`s

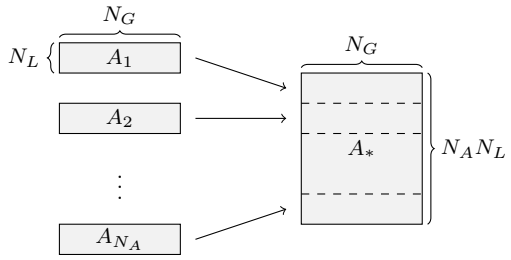

Figure 2: Memory Layout of A_a in A_*

This algorithm performs a total of $4N_A N_L N_G^2$ FLOPs[11] in N_A calls to `zherk`. To obtain library invocations with potentially higher performance, we improve upon Listing 1 and combine all N_A `zherk` invocations into a single, large one by choosing a smart data layout for the A_as; we stack all A_a vertically on top of each other into a single matrix A_* as shown in Figure 2. Here, A_* is of size $N_A N_L \times N_G$ and has leading dimension $N_A N_L$. With this memory layout for A_*, Listing 1 turns into a single call to `zherk`.

```
1  S += A_*^H A_*                                   (zherk: 4 N_A N_L N_G^2 FLOPs)
```

Listing 2: $S \mathrel{+}= S_{AA}$ by one `zherk`

4.3. Computing S: Adding S_{BB}

The term S_{BB} in S is very similar to S_{AA} (Eq. (17)):

$$S \mathrel{+}= S_{BB} = \sum_{a=1}^{N_A} B_a^H \dot{U}_a^H \dot{U}_a B_a \ . \tag{18}$$

First, we logically distribute the $\dot{U}_a$ symmetrically to B_a^H and B_a as $B'_a \coloneqq \dot{U}_a B_a$, allowing us to rewrite Eq. (18) as

$$S \mathrel{+}= S_{BB} = \sum_{a=1}^{N_A} {B'_a}^H B'_a \ .$$

At this point, by applying the same memory layout to the B_a (packing them into B_*), the entire $S \mathrel{+}= S_{AA} + S_{BB}$ is computed in the following Listing.

```
1  S += A_*^H A_*                                   (zherk: 4 N_A N_L N_G^2 FLOPs)
2  B_* := Ů_* B_*                                   (2 N_A N_L N_G FLOPs)
3  S += B_*^H B_*                                   (zherk: 4 N_A N_L N_G^2 FLOPs)
```

Listing 3: $S \mathrel{+}= S_{AA} + S_{BB}$ (final)

Here, line 2 scales the rows of B_*: $\dot{U}_*$ is diagonal and represents the concatenation of the $\dot{U}_a$ across all atoms a.

[11] N_A = number of atoms, N_L = number of spherical harmonics, and N_G = number of basis functions.

4.4. Computing H: Constructing H_{AB+BA}

We begin the computation of H by constructing the component

$$H \mathrel{+}= H_{AB+BA} = \sum_{a=1}^{N_A} A_a^H T_a^{[AB]} B_a + B_a^H T_a^{[BA]} A_a \ . \tag{19}$$

Since $T_a^{[BA]} = \left(T_a^{[AB]}\right)^H$, this term is Hermitian and can be rewritten as follows:

$$H_{AB+BA} = \sum_{a=1}^{N_A} \left(T_a^{[BA]} A_a\right)^H B_a + B_a^H \left(T_a^{[BA]} A_a\right) \ .$$

Introducing $X_a \coloneqq T_a^{[BA]} A_a$ as an intermediate, we obtain

$$H_{AB+BA} = \sum_{a=1}^{N_A} X_a^H B_a + B_a^H X_a \ , \tag{20}$$

matching the BLAS kernel `zher2k`, which computes only the lower triangular half of the symmetric H_{AB+BA} (including the diagonal). This leads to the following algorithm:

```
1  for a ≔ 1, ..., N_A:
2    X_a ≔ T_a^[BA] A_a                    (zgemm: 8 N_L^2 N_G FLOPs)
3    H += X_a^H B_a + B_a^H X_a            (zher2k: 8 N_L N_G^2 FLOPs)
```

Listing 4: $H \mathrel{+}= H_{AB+BA}$ with N_A `zher2k`s

While we cannot combine the `zgemm`s into a single kernel, by applying the memory layout used in A_* and B_* to the X_a, stacking them in X_*, we can replace the N_A calls to `zher2k` by a single one:

```
1  for a ≔ 1, ..., N_A:
2    X_a ≔ T_a^[BA] A_a                    (zgemm: 8 N_L^2 N_G FLOPs)
3    add X_a to X_*
4    add B_a to B_*
5  H += X_*^H B_* + B_*^H X_*              (zher2k: 8 N_A N_L N_G^2 FLOPs)
```

Listing 5: $H \mathrel{+}= H_{AB+BA}$

This algorithm performs a total of $8N_A N_L^2 N_G + 8N_A N_L N_G^2$ FLOPs.

4.5. Computing H: Incorporating H_{BB}

Next we consider the contribution

$$H \mathrel{+}= H_{BB} = \sum_{a=1}^{N_A} B_a^H T_a^{[BB]} B_a \ . \tag{21}$$

Since $T_a^{[BB]} = \left(T_a^{[BB]}\right)^H$ is Hermitian, so is this entire term. Using this property, we can carefully rewrite H_{BB} in a similar way as we did for H_{AB+BA} in Eq. (19) above:

$$\begin{aligned} H_{BB} &= \sum_{a=1}^{N_A} B_a^H T_a^{[BB]} B_a \\ &= \sum_{a=1}^{N_A} \frac{1}{2} B_a^H \left(T_a^{[BB]}\right)^H B_a + \frac{1}{2} B_a^H T_a^{[BB]} B_a \\ &= \sum_{a=1}^{N_A} \left(\tfrac{1}{2} T_a^{[BB]} B_a\right)^H B_a + B_a^H \left(\tfrac{1}{2} T_a^{[BB]} B_a\right) \end{aligned}$$

Now, introducing the intermediate $Y_a \coloneqq \frac{1}{2} T_a^{[BB]} B_a$, we arrive at

$$H_{BB} = \sum_{a=1}^{N_A} Y_a^H B_a + B_a^H Y_a \ .$$

Noting the similarity with Eq. (20), we combine $Z_a \coloneqq X_a + Y_a$ to compute $H_{AB+BA} + H_{BB}$ in a single `zher2k` of unchanged size. Using the same memory layout for the Z_a as for the X_a before, i.e., stacking them into Z_*, we arrive at the following algorithm:

```
1  for a := 1, ..., N_A:
2    Z_a := T_a^[BA] A_a                          (zgemm: 8 N_L^2 N_G FLOPs)
3    Z_a += 1/2 T_a^[BB] B_a                      (zhemm: 8 N_L^2 N_G FLOPs)
4    add Z_a to Z_*
5    add B_a to B_*
6  H += Z_*^H B_* + B_*^H Z_*                     (zher2k: 8 N_A N_L N_G^2 FLOPs)
```

Listing 6: $H \mathrel{+}= H_{AB+BA} + H_{BB}$ (final)

This algorithm performs $16 N_A N_L^2 N_G + 8 N_A N_L N_G^2$ FLOPs, which means that integrating the contribution into Eq. (21) only costs $8 N_A N_L^2 N_G$ FLOPs.

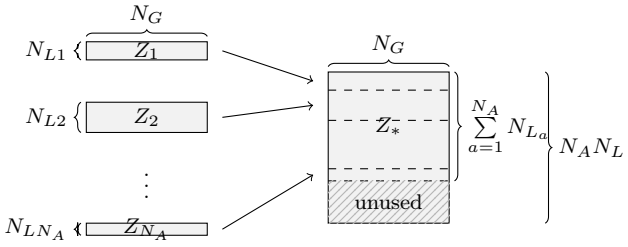

Figure 3: Memory layout of $Z_a \in \mathbb{C}^{N_{L_a} \times N_G}$ in Z_* with varying N_{L_a} for Listing 6

4.6. Computing H: Updating H += H_{AA}

The only term remaining to construct H is

$$H \mathrel{+}= H_{AA} = \sum_{a=1}^{N_A} A_a^H T_a^{[AA]} A_a \ . \tag{22}$$

Unfortunately, although very similar to H_{BB} in Eq. (21), this update cannot be fused into the `zher2k` along with both H_{AA} and H_{AB+BA} since there is generally no factor common to all three of terms H_{AA}, H_{AB+BA}, and H_{BB}: we can either compute $H_{AB+BA} + H_{BB}$ together by factoring out B^* (see subsection 4.5 and Listing 6) or apply the same approach to $H_{AA} + H_{AB+BA}$ using the common A^*, yet not both simultaneously.

We chose to fuse H_{BB} into H_{AB+BA} and now have to deal with the leftover H_{AA}. Our method for computing the contribution H_{AA} depends on the properties of the $T_a^{[AA]}$:

- In the general case, we can not exploit the Hermitian symmetry of H_{AA}, and are forced to use non-Hermitian kernels which perform some redundant computation. This is discussed in subsubsection 4.6.1.
- If the $T_a^{[AA]}$ are Hermitian positive definite (HPD), we can retain the Hermitian-ness by Cholesky decomposing and evenly distributing the $T_a^{[AA]}$. This is discussed in subsubsection 4.6.2.

Since in practice we commonly encounter both HPD and non-HPD $T_a^{[AA]}$ in the same system, in subsubsection 4.6.3, we combine both approaches dynamically, depending on HPD-ness of the individual $T_a^{[AA]}$.

4.6.1. The General Case

When the $T_a^{[AA]}$ are not HPD, we cannot turn $A_a^H T_a^{[AA]} A_a$ into a symmetric series of library calls. Hence we are left with computing $X_a \coloneqq T_a^{[AA]} A_a$ with a call to `zhemm` and then summing

$$H \mathrel{+}= H_{AA} = \sum_{a=1}^{N_A} A_a^H X_a \ .$$

Once more using the established memory layout of A_* and X_*, this leads to the following algorithm to update $H \mathrel{+}= H_{AA}$:

```
1  for a ≔ 1, ..., N_A:
2    X_a ≔ T_a^[AA] A_a                    (zhemm: 8 N_L^2 N_G FLOPs)
3    add X_a to X_*
4    add A_a to A_*
5  H += A_*^H X_*                          (zgemm: 8 N_A N_L N_G^2 FLOPs)
```

Listing 7: $H \mathrel{+}= H_{AA}$ for exclusively non-HPD $T_a^{[AA]}$

Note that this algorithm performs $8N_A N_L^2 N_G + 8N_A N_L N_G^2$ FLOPs to add H_{AA} to H, while Listing 6 incorporated H_{BB} into the computation of H_{AB+BA} with only $8N_A N_L^2 N_G$ FLOPs.

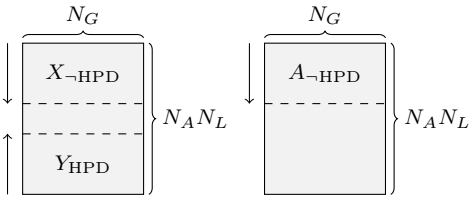

Figure 4: Memory layout and stacking direction of $X_{\neg\text{HPD}}$, Y_{HPD}, and $A_{\neg\text{HPD}}$ in two matrices.

4.6.2. The Hermitian Positive Definite Case

If the $T_a^{[AA]}$ are HPD, we can Cholesky decompose them as $C_a C_a^H \coloneqq T_a^{[AA]}$, where C_a is lower triangular. This allows us to rewrite H_{AA} as follows:

$$H_{AA} = \sum_{a=1}^{N_A} A_a^H T_a^{[AA]} A_a = \sum_{a=1}^{N_A} A_a^H C_a C_a^H A_a = \sum_{a=1}^{N_A} \left(C_a^H A_a\right)^H \left(C_a^H A_a\right) \ .$$

Defining the intermediate $Y_a \coloneqq C_a^H A_a$, we can rewrite this into the symmetric update

$$H \mathrel{+}= H_{AA} = \sum_{a=1}^{N_A} Y_a^H Y_a \ .$$

Applying our memory layout for Y_*, we arrive at the following algorithm:

```
1  for a ≔ 1, ..., N_A:
2    C_a ≔ Chol(T_a^[AA])                    (zpotrf: 4/3 N_L^3 + O(N_L^2) FLOPs)
3    Y_a ≔ C_a^H A_a                         (ztrmm: 4 N_L^2 N_G FLOPs)
4    add Y_a to Y_*
5  H += Y_*^H Y_*                            (zherk: 4 N_A N_L N_G^2 FLOPs)
```

Listing 8: $H \mathrel{+}= H_{AA}$ for exclusively HPD $T_a^{[AA]}$

This algorithm performs $\frac{4}{3}N_A N_L^3 + O(N_A N_L^2) + 4N_A N_L^2 N_G + 4N_A N_L N_G^2$ FLOPs, which, neglecting the lower order contribution of `zpotrf` (since $N_L \ll N_G$), performs only half as many flops as Listing 7 for the non-HPD case.

4.6.3. Dynamic Testing for HPD-ness

In practice, it is common to encounter a mixture of HPD and non-HPD $T_a^{[AA]}$ in the same system. To still benefit from the cheaper algorithm for the HPD case, we combine Listing 7 and Listing 8 and dynamically decide which route to take depending on the HPD-ness of each separate $T_a^{[AA]}$. As a result the majority of the computation is performed by not one but two large updates (`zgemm` and `zherk`) outside the loop over the atoms a.

We check whether each $T_a^{[AA]}$ is HPD by attempting to Cholesky decompose them: if the decomposition succeeds, we continue as in Listing 8; if it fails (which `zpotrf` indicates via an error code), we follow the general Listing 7. This is summarized in the following algorithm:

```
1  for a ≔ 1, ..., N_A:
2    try:
3      C_a ≔ Chol(T_a^[AA])                  (zpotrf: 4/3 N_L^3 + O(N_L^2) FLOPs)
4    success:
```

```
5     Y_a ≔ C_a^H A_a                                   (ztrmm: 4N_L^2 N_G FLOPs)
6   failure:
7     X_a ≔ T_a^[AA] A_a                                (zhemm: 8N_L^2 N_G FLOPs)
```

Listing 9: Dynamic branching based on the HPD-ness of $T_a^{[AA]}$

Note that this algorithm is missing the actual update to compute $H \mathrel{+}= H_{AA}$. Since these are different in Listing 7 and Listing 8, we have to modify the memory layout of our matrices to still perform them with a single call to each `zgemm` and `zherk`: for the cases where $T_a^{[AA]}$ is HPD, we collect the Y_a in a matrix Y_{HPD}; for the non-HPD $T_a^{[AA]}$, we collect the X_a and the corresponding A_a in, respectively, $X_{\neg\text{HPD}}$ and $A_{\neg\text{HPD}}$. To minimize the memory impact, we store Y_{HPD} and $X_{\neg\text{HPD}}$ in a single matrix of size $N_A N_L \times N_G$ by stacking the Y_a from the bottom, while stacking X_a from the top. $A_{\neg\text{HPD}}$ is simultaneously stacked from the top in the original A_* matrix, potentially overwriting already processed A_a, as shown in Figure 4. With this memory layout, we arrive at the following algorithm:

```
1   for i += 1, ..., N_A:
2     try:
3       C_a ≔ Chol(T_a^[AA])                            (zpotrf: 4/3 N_L^3 + O(N_L^2) FLOPs)
4     success:
5       Y_a ≔ C_a^H A_a                                 (ztrmm: 4N_L^2 N_G FLOPs)
6       add Y_a to Y_HPD
7     failure:
8       X_a ≔ T_a^[AA] A_a                              (zhemm: 8N_L^2 N_G FLOPs)
9       add X_a to X_¬HPD
10      add A_a to A_¬HPD
11  H += A_¬HPD^H X_¬HPD                                (zgemm: 8N_{A_¬HPD} N_L N_G^2 FLOPs)
12  H += Y_HPD^H Y_HPD                                  (zherk: 4N_{A_HPD} N_L N_G^2 FLOPs)
```

Listing 10: $H \mathrel{+}= H_{AA}$ (final)

4.7. Accounting for Differently Sized $T_a^{[\dots]}$

In the previous sections, to simplify the discussion of our algorithms, we assumed that all $T_a^{[\dots]}$ are of size $N_L \times N_L$. In practice however, for different atoms a, $T_a^{[\dots]}$ can be of a different size $N_{L_a} \times N_{L_a}$, where $N_{L_a} \leq N_L$. The simplest way to deal with these differing sizes would be to pad the matrices with zero rows and columns up to $N_L \times N_L$. However, doing so would lead to numerous redundant computations on zero entries in the computation of H. Instead, we perform only the necessary work by omitting the zero rows in the layout of the composite matrices (marked with the subscripts "$_*$", "$_{\neg\text{HPD}}$", and "$_{\text{HPD}}$") thereby compacting them for the large updates on H throughout $H \mathrel{+}= H_{AB+BA} + H_{BB}$ (Listing 6) and $H \mathrel{+}= H_{AA}$ (Listing 10).

In Listing 6, we have $T_a^{[BA]}, T_a^{[BB]} \in \mathbb{C}^{N_{L_a} \times N_{L_a}}$ with $N_{L_a} \leq N_L$. These matrices will only multiply the first N_{L_a} rows of A_a and B_a in, respectively, lines 2 and 3 of Listing 6. The resulting intermediate $Z_a = T_a^{[BA]} A_a + \frac{1}{2} T_a^{[BB]} B_a$ is therefore of size $N_{L_a} \times N_G$. As before, we collect these Z_a, now of varying sizes, one below the other in Z_*. However, to perform the large `zher2k` $H \mathrel{+}= Z_*^H B_*$ in line 6 of Listing 6, we need to pack B_* accordingly. We achieve such a result by only keeping the relevant N_{L_a} rows of each B_a in B_*, aligning them with the corresponding rows of Z_a in Z_*. After this process, the Z_* (and B_* analogously) are in the memory layout shown in Figure 3 and $H \mathrel{+}= Z_*^H B_*$ can be performed as before in Listing 6.

We apply the same concept to Listing 10, where $T_a^{[AA]} \in \mathbb{C}^{N_{L_a} \times N_{L_a}}$ with $N_{L_a} \leq L$. However, due to the separation of the updates according to the HPD-ness of these $T_a^{[AA]}$, the resulting memory layout is slightly different. For the HPD cases, the $X_a \in \mathbb{C}^{N_{L_a} \times N_G}$ and the corresponding first N_{L_a} rows of A_a are collected in, respectively, $X_{\neg\text{HPD}}$ and $A_{\neg\text{HPD}}$ just as Z_* and B_* above. For the non-HPD cases, the $Y_a \in \mathbb{C}^{N_{L_a} \times N_G}$ are collected contiguously in Y_{HPD}. As before, by stacking the Y_a from the bottom, $X_{\neg\text{HPD}}$ and Y_{HPD} share the same buffer. Altogether, these modifications of the layout allow us to perform the large updates $H \mathrel{+}= A_{\neg\text{HPD}}^H X_{\neg\text{HPD}}$ and $H \mathrel{+}= Y_{\text{HPD}}^H Y_{\text{HPD}}$ in the same fashion as detailed in Listing 10.

4.8. Composing the Complete Algorithm

The computation of both S and H is entirely contained in the Listings 3, 6, and 10 which can be thought of as the main components. The complete algorithm is constructed by selecting an order of execution for these three components while considering their memory footprint. We start by analyzing the memory requirements of each Listing and which are the matrices overwritten in the course of their execution. In the following, all matrices with relevant memory impact (of size $N_A N_L \times N_G$ or $N_G \times N_G$) are highlighted:

- Listing 3 computes $S \mathrel{+}= S_{AA} + S_{BB} = \sum_{a=1}^{N_A} A_a^H A_a + B_a^H \dot{U}_a^H \dot{U}_a B_a$

  Input: $S \in \mathbb{C}^{N_G \times N_G}$, $A_* \in \mathbb{C}^{N_A N_L \times N_G}$, $B_* \in \mathbb{C}^{N_A N_L \times N_G}$

  Temporaries: $(\dot{U}_* B_*) \in \mathbb{C}^{N_A N_L \times N_G}$

- Listing 6 computes $H \mathrel{+}= H_{AB+BA} + H_{BB} = \sum_{a=1}^{N_A} A_a^H T_a^{[AB]} B_a + B_a^H T_a^{[BA]} A_a + B_a^H T_a^{[BB]} B_a$

  Input: $H \in \mathbb{C}^{N_G \times N_G}$, $A_* \in \mathbb{C}^{N_A N_L \times N_G}$, $B_* \in \mathbb{C}^{N_A N_L \times N_G}$, $N_A \times T_a^{[BA]} \in \mathbb{C}^{N_L \times N_L}$, $N_A \times T_a^{[BB]} \in \mathbb{C}^{N_L \times N_L}$

  Temporaries: $Z_* \in \mathbb{C}^{N_A N_L \times N_G}$, $B_* \in \mathbb{C}^{N_A N_L \times N_G}$ (with new memory layout, Figure 3)

- Listing 10 computes $H \mathrel{+}= H_{AA} = \sum_{a=1}^{N_A} A_a^H T_a^{[AA]} A_a$

  Input: $H \in \mathbb{C}^{N_G \times N_G}$, $A_* \in \mathbb{C}^{N_A N_L \times N_G}$, $N_A \times T_a^{[AA]} \in \mathbb{C}^{N_L \times N_L}$

  Temporaries: $C_a \in \mathbb{C}^{N_L \times N_L}$, $A_{\neg\text{HPD}} \in \mathbb{C}^{N_A N_L \times N_G}$, $(X_{\neg\text{HPD}}, Y_{\text{HPD}}) \in \mathbb{C}^{N_A N_L \times N_G}$

In summary, all three components work on one matrix of size $N_G \times N_G$ (either S or H) and use two large buffers of size $N_A N_L \times N_G$. Since in the end both S and H must coexist in memory, this means we need a minimum of $32 N_G^2$ bytes for H and S (2 double precision floats of 8 bytes per number) and $32 N_A N_L N_G$ bytes for the two other buffers.

Both Listing 3 and Listing 6 require both A_* and B_* as inputs, while Listing 10 only requires A_*; at the same time, both Listing 6 and Listing 10 overwrites both A_* and B_*, while Listing 3 only overwrites B_*. As a result, we either need to create A_* and B_* at least twice, or keep copies of them in memory. We minimize this overhead by arranging the algorithms as follows, where the flag *backup* indicates whether to keep copies of A_* and B_* or to re-create them:

1	create A_* and B_*	$+32N_AN_LN_G$ bytes
2	**if** *backup*:	
3	back up $A'_* \coloneqq A_*$ and $B'_* \coloneqq B_*$	$+32N_AN_LN_G$ bytes
4	$H \coloneqq 0 \in \mathbb{C}^{N_G \times N_G}$	$+16N_G^2$ bytes
5	$H \mathrel{+}= H_{AB+BA} + H_{BB}$	(Listing 6)
6	**if** *backup*:	
7	restore $A_* \coloneqq A'_*$ and $B_* \coloneqq B'_*$	
8	free A'_* and B'_*	$-32N_AN_LN_G$ bytes
9	**else**:	
10	create A_* and B_*	
11	$S \coloneqq 0 \in \mathbb{C}^{N_G \times N_G}$	$+16N_G^2$ bytes
12	$S \mathrel{+}= S_{AA} + S_{BB}$	(Listing 3)
13	$H \mathrel{+}= H_{AA}$	(Listing 10)

Listing 11: Final Algorithm

On the right of Listing 11, we indicate the large memory (de-)allocations preceded by the (minus)plus sign. Depending on the *backup* flag, the total memory requirement of the algorithm assumes different values:

- *backup* = **true**: Up to line 5, we allocated $64N_AN_LN_G + 32N_G^2$ bytes, then deallocate $32N_AN_LN_G$ bytes in line 8 and allocate another $16N_G^2$ bytes in line 11. As a result, this scenario overall requires $32N_AN_LN_G + 16N_G^2 + \max(32N_AN_LN_G, 16N_G^2)$ bytes, which can be expressed as

$$32N_AN_LN_G + 32N_G^2 + \max(32N_AN_LN_G - 16N_G^2, 0) \text{ bytes.}$$

- *backup* = **false**: This algorithm allocates a total of

$$32N_AN_LN_G + 32N_G^2 \text{ bytes.}$$

Although the creation of A_* and B_* is fast compared to the computations performed in our algorithms, experiments have shown that they can account for about 4% of the total compute time. Keeping copies of these matrices on the other hand is negligible in terms of time overhead. In conclusion, we recommend to use the back-up mechanism unless $2N_AN_L > N_G$ and the used machine poses a practically relevant memory limitation.

Example. To provide a feeling for both the memory footprint and where HSDLA spends its compute time in practice, we consider two examples: One of the numerical tests we run involved a physical system with $N_A = 512$, $N_L = 49$, and $N_G = 2256$, which requires

$$32 \cdot 512 \cdot 49 \cdot 2256 + 32 \cdot 2256^2 + \max(32 \cdot 512 \cdot 49 \cdot 2256 - 16 \cdot 2256^2, 0) \text{ bytes} = 3.45\,\text{GiB},$$

and can still fit in the memory of modern laptops. In this example, 98.06 % of the computation performed by HSDLA is covered by the large updates of H and S (`zherk`s in Listing 3, `zher2k` in Listing 6, and `zgemm` / `zherk`[12] in Listing 10).

[12] Assuming 50 % of the $T_a^{[AA]}$ are HPD.

Increasing $\mathbf{K}_{\max}$ for improved accuracy raises N_G to 7946 and expands the memory footprint to 18.5 GiB. While such a problem comfortably fits in the main memory of common cluster nodes, other practical configurations, such as $N_A = 384$, $N_L = 81$ and $N_G = 29\,144$ with a footprint of 66.7 GiB require larger shared memory systems. The portion of the computations spent on the updates of H and S is increased to 99.75 %.

5. Performance results

In this section we present performance tests for two distinct atomic systems, colloquially referred to as NaCl and TiO_2. By including both a conductor and an insulator, these systems represent a heterogeneous sample with different physical properties. All tests were performed by timing the generation of **H** and **S** for only a single **k**-point. We repeated the simulations varying $\mathbf{K}_{\max}$ over a range of values going from 2.5 to 4.0, in steps of size 0.1. Since the value of $\mathbf{K}_{\max}$ determines the size N_G of the basis set, by increasing the value of this parameter one can simulate the atomic systems with greater accuracy. This improved accuracy comes at the cost of a larger memory footprint and additional computations, resulting in longer execution times.

We compare the execution times for our implementation HSDLA with those for FLEUR v.26e, and use two hardware platforms: the first platform consists of two 10 core IvyBridge-EP E5-2680 v2 with 256 GiB of main memory; the second platform consists of two 12 core Haswell-EP E5-2680 v3 with 64 GiB of main memory. Our algorithm uses the platforms' parallelism through the BLAS and LAPACK routines from Intel's Math Kernel Library (MKL) version 11.3. On the other hand, FLEUR, which is linked to the same MKL library, is parallelized using MPI and linked to IntelMPI version 5.0.

Figure 5 presents the execution times for both the NaCL system (top) and the TiO_2 system (bottom), with increasing $\mathbf{K}_{\max}$. The timings are obtained on both the IvyBridge platform (——/——) and the Haswell platform (——/——) using either a single core (left), or all of the platforms cores (right). Note that on the Haswell platform timings for the TiO_2 system are only available up to $\mathbf{K}_{\max} = 3.6$, since beyond this point the main memory capacity of 64 GiB is exceeded.

First, we observe that HSDLA (——/——) is consistently faster than FLEUR (——/——) across all setups. With multi-threading (right panel), the performance fluctuates considerably more for FLEUR than for HSDLA, yet the trend remains the same with HSDLA yielding significant performance improvements over FLEUR.

	$\mathbf{K}_{\max} = 2.5$			$\mathbf{K}_{\max} = 3.0$			$\mathbf{K}_{\max} = 3.5$			$\mathbf{K}_{\max} = 4.0$		
	HSDLA	FLEUR	×	HSDLA	FLEUR	×	HSDLA	FLEUR	×	HSDLA	FLEUR	×
NaCl	14.24	13.61	**1.05**	29.99	42.11	**1.40**	63.89	104.11	**1.66**	116.37	242.97	**2.09**
TiO_2	79.95	146.48	**1.83**	189.54	398.49	**2.10**	382.44	967.77	**2.53**			

Table 1: Runtime in seconds and speedups (in **bold**) of HSDLA over FLEUR with increasing $\mathbf{K}_{\max}$ on all 24 cores of the Haswell.

Recall that the size N_G of **H** and **S** is approximately equal to the product of a prefactor (50-80) times the number of atoms N_A. The greater the $\mathbf{K}_{\max}$, the larger is the used prefactor – typically a $\mathbf{K}_{\max} = 4.0$ implies a prefactor equal to 80. At first glance, the plots in Figure 5

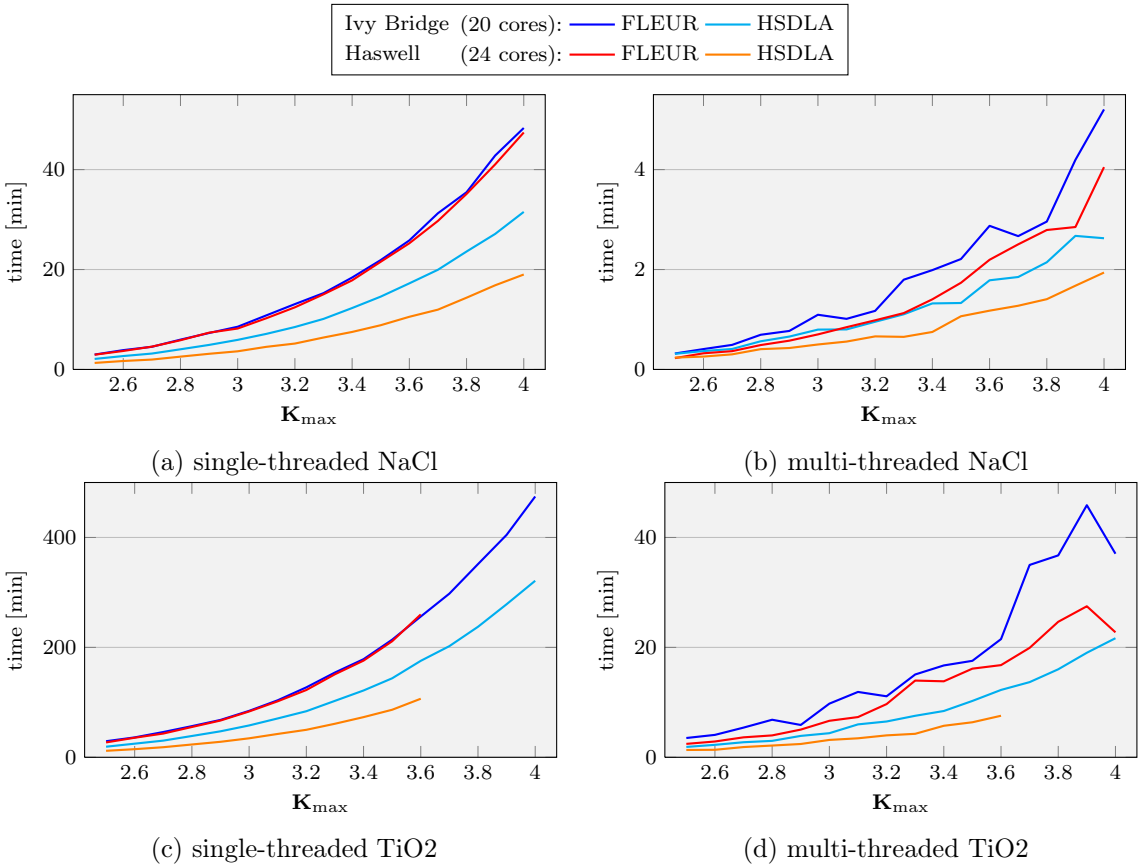

(a) single-threaded NaCl

(b) multi-threaded NaCl

(c) single-threaded TiO2

(d) multi-threaded TiO2

Figure 5: Execution time of FLEUR and HSDLA for different setups and increasing $\mathbf{K}_{\max}$

hint that higher values of $\mathbf{K}_{\max}$ favor the use of HSDLA over FLEUR in terms of execution time. Such a hint is confirmed by the data in Table 2, where speedup values are shown to get larger for increasing values of $\mathbf{K}_{\max}$. This observation indirectly suggests that simulations with a larger number of atoms – which increase N_A instead of the prefactor – would also be favored by HSDLA. Since a larger N_A implies larger matrices in Listing 11, the conclusion just drawn confirms what is the conventional wisdom when dealing with level 3 BLAS kernels: the larger the matrices, the greater the performance of the basic linear algebra kernels.

In Figure 6 we present the speedup of HSDLA over FLEUR with increasing number of threads. In the single-threaded case (left) we obtain consistent speedups of 1.5× on IvyBridge and 2.4× on Haswell. This result shows that HSDLA not only makes better use of the available resources but is also performance portable: in contrast to FLEUR, which takes the same time on the older IvyBridge architecture as on the more modern Haswell, HSDLA makes better use of the newer processor's increased performance.

A small digression is in order here to understand why the speedup of HSDLA relative to FLEUR is 2× and not higher. As already mentioned earlier, by maintaining the dense linear algebra structure, HSDLA can leverage high-performance kernels like BLAS and LAPACK. It is well understood that on a single thread, a BLAS Level 3 kernel executing the same number of

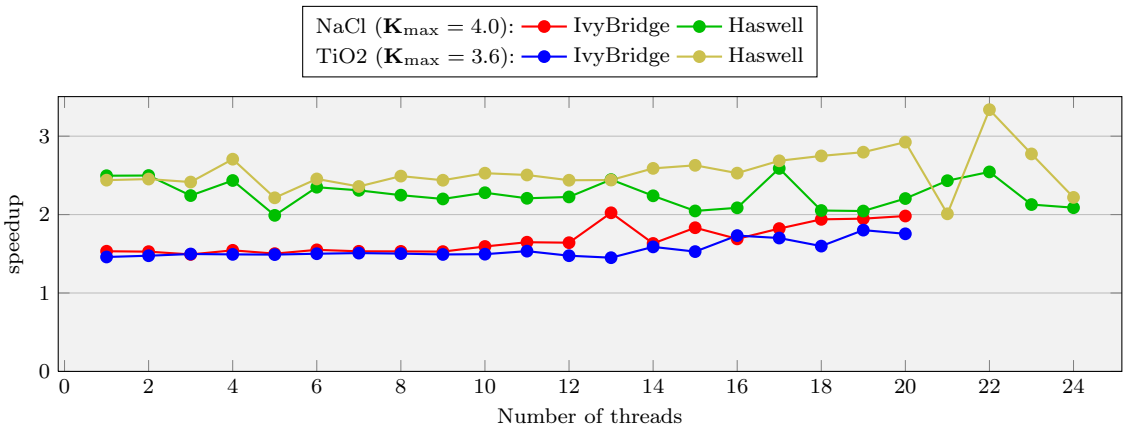

Figure 6: Speedup of our algorithm over FLEUR with increasing parallelism

	NaCl ($\mathbf{K}_{\max} = 4.0$)						TiO_2 ($\mathbf{K}_{\max} = 3.6$)					
	IvyBridge			Haswell			IvyBridge			Haswell		
	HSDLA	FLEUR	×	HSDLA	FLEUR	×	HSDLA	FLEUR	×	HSDLA	FLEUR	×
1 core	31.53	48.31	**1.53**	19.00	47.41	**2.50**	175.53	256.15	**1.46**	106.56	259.91	**2.44**
2 cores	16.10	24.58	**1.53**	9.98	24.95	**2.50**	86.68	127.90	**1.48**	53.48	131.21	**2.45**
1 CPU	3.90	6.21	**1.59**	2.25	5.00	**2.22**	19.63	29.35	**1.50**	10.63	25.95	**2.44**
2 CPUs	2.61	5.20	**1.99**	1.93	4.03	**2.09**	12.25	21.50	**1.76**	7.55	16.76	**2.22**

Table 2: Scalability of HSDLA and FLEUR: execution times in minutes on Haswell (12 cores / CPU) and IvyBridge (10 cores / CPU); speedups of HSDLA over FLEUR in **bold**.

FLOPs can be up to one order of magnitude faster[13] than a nested loop. If HSDLA and FLEUR were executing the same number of FLOPs, one would have expected a speedup quite larger than the one we measure. On the other hand, as we mention in Sec. 3, FLEUR takes advantage of mathematical simplifications which minimize significantly the number of executed FLOPs with respect to the linear algebra structure of the original mathematical formulation. Since the difference in FLOP count between FLEUR and HSDLA is roughly two orders of magnitude for the system considered here, one could expect to obtain a very limited or no speedup at all. In practice, the significantly larger number of FLOPs executed reduces the speedup that HSDLA could achieve but does not cancel out the advantage of using BLAS Level 3 kernels. The net result is that HSDLA still clearly outperforms FLEUR. In an extension of the present work, one could investigate the feasibility of including, whenever possible, some of FLEUR's mathematical simplifications into HSDLA and increasing the overall speedup.

Table 2 presents the execution times of HSDLA and FLEUR for varying numbers of threads. Both FLEUR and HSDLA scale very well: on both architectures, the parallel efficiency on 1 CPU is almost 80 % while on 2 CPUs it is still around 60 %. From Table 2 it is evident that on the shared-memory architecture, FLEUR and HSDLA have similar scalability and hence similar parallel efficiency. Note that the comparison is between an MPI parallelization (FLEUR) and a multi-threaded one (HSDLA uses MKL). Despite the clear differences between

[13] On multi-threaded platforms the advantage given by BLAS can be even higher.

these two standards, their performance signature on shared-memory platforms are comparable. Stepping out of the single computing node and moving into the distributed memory arena, we expect that the performance signature will change dramatically. In particular, since most of the generation of **H** and **S** in FLEUR is carried out in nested loops, we expect a drop in parallel efficiency with respect to any optimized implementation of a distributed BLAS library [23].

Overall, these performance results suggest that by restructuring FLEUR according to the HSDLA algorithm, one significantly improves its performance and achieves performance portability.

6. Summary and conclusions

Extending the life span and increasing the functionalities of legacy codes like FLEUR has become a central issue in many areas of computational science and engineering. Enabling such codes to access and use efficiently massively parallel architectures is key. In this context, performance portability plays a vital role in empowering legacy codes to adapt to the emergence of heterogeneous computing platforms.

In the present work we focus on FLEUR, a Density Functional Theory code based on the FLAPW method and developed at the Forschungszentrum Jülich over the course of more than two decades. Specifically, we focus on the generation of the Hamiltonian and Overlap matrices, one of the most computationally intensive sections of the code. Together with their eigendecomposition, the generation of these matrices accounts for more than 80 % of the total CPU time.

Similar to many other scientific codes, FLEUR has been implemented by translating the mathematical formulation directly into code. In addition, optimizations aimed at reducing the floating point operations and the memory footprint were included early on, making it difficult to incorporate later extensions. With the advent of massively parallel computing architectures, FLEUR has the opportunity to simulate atomic systems with an unprecedented number of atoms. Unfortunately, the inherently cache-memory insensitive structure prevents FLEUR from taking advantage of modern hardware architectures with the consequent loss in performance and portability.

In order to lift these limitations and create a performance portable code, we abstracted from the current implementation and went back to its mathematical foundation. First, we expressed the fundamental expressions involved in the creation of the **H** and **S** matrices in terms of a combination of dense linear algebra operations. Then, we cast these operations as kernels supported by well established high-performance libraries. In addition, we applied a number of high-level optimizations by combining operations and by reducing the number of redundant computations. The resulting algorithm, HSDLA, is an efficient and portable implementation that attains high-performance on shared memory architectures. Compared to the original FLEUR implementation, HSDLA achieves speedups between 1.5× and 2.5×. Most importantly, when tested on newer architectures, HSDLA increases the speedup gap over FLEUR, showing its increased performance portability.

Building on these promising initial results, we intend to investigate how FLEUR's mathematical FLOP count reductions (see Sec. 3) can be applied to the high-level linear algebra formulation, and incorporated while maintaining our algorithm's modularity and BLAS-based performance portability. Furthermore, thanks to the modularity and the high-level optimizations used, extending the HSDLA algorithm to distributed memory platforms would only

require to adapt the implementation to high-performance *distributed* dense linear algebra libraries (e.g. Elemental [23]). From such an extension, we expect to obtain an increased scalability and achieve an enhanced parallelism with little effort.

Appendix A. Some details of the FLAPW method

This section provides supplementary material to Sec. 2, and is meant to fill some of the gaps that were inevitably left open for the sake of conciveness. While this appendix clarifies the mathematical setup of the FLAPW method, it is by no means necessary to understand and capture the message of this paper.

Appendix A.1. Definition of the A and B tensors

Introduced in Eq. (8) and at the center of the present work, the A and B constant coefficients are explicitly determined by differentiating and matching in value both parts of Eq. (8) at the muffin tin/interstitial boundary,[14]

$$A_{(l,m),a,t}(\mathbf{k}) = \frac{4\pi}{W_{l,a}\sqrt{\Omega}} i^l \exp\left(i\mathbf{K}_t \cdot \mathbf{x}_a\right) Y^*_{l,m}\left(\mathbf{R}_a \hat{K}_t\right) \cdot \\ \cdot \left[\dot{u}_{l,a}\left(r_{MT,a}\right) K_t \, j'_l\left(r_{MT,a}\, K_t\right) - \dot{u}'_{l,a}\left(r_{MT,a}\right) j_l\left(r_{MT,a}\, K_t\right)\right] \quad \text{(A.1)}$$

and

$$B_{(l,m),a,t}(\mathbf{k}) = \frac{4\pi}{W_{l,a}\sqrt{\Omega}} i^l \exp\left(i\mathbf{K}_t \cdot \mathbf{x}_a\right) Y^*_{l,m}\left(\mathbf{R}_a \hat{K}_t\right) \cdot \\ \cdot \left[-u_{l,a}\left(r_{MT,a}\right) K_t \, j'_l\left(r_{MT,a}\, K_t\right) + u'_{l,a}\left(r_{MT,a}\right) j_l\left(r_{MT,a}\, K_t\right)\right] \quad \text{(A.2)}$$

with the Wronskian defined as

$$W_{l,a} = \dot{u}_{l,a}\left(r_{MT,a}\right) u'_{l,a}\left(r_{MT,a}\right) - u_{l,a}\left(r_{MT,a}\right) \dot{u}'_{l,a}\left(r_{MT,a}\right).$$

The functions j_l and j'_l are the spherical Bessel functions (and their derivatives) which are the result of the Rayleigh expansion of plane waves $\exp\left(i\mathbf{K}_t \cdot \mathbf{r}\right)$ in spherical harmonics. Notice that a solution $u_{l,a}$ to the radial equation (7) is required for all values of l and each atom positioned at $\mathbf{x}_a$. In most cases, there are multiple atoms at different positions that share the same muffin-tin geometry (e.g. of the same element in the periodic table), so that the radial equation need to be solved only once per atom species. In doing so, a rotation matrix needs to be introduced in the spherical harmonics $\mathbf{R}_a$ to account for atoms of the same type at distinct positions.

The different derivatives of $u'_{l,a}$ and $\dot{u}'_{l,a}$ result from the differentiation of Eq. (8) by its radial argument, which occurs when we require

$$\partial_{\mathbf{r}}\, \varphi\left(\mathbf{r}\right)\Big|_{\mathrm{INT}} \equiv \partial_{\mathbf{r}}\, \varphi\left(\mathbf{r}\right)\Big|_{\mathrm{MT}}.$$

[14] For a detailed description, see [16, pp. 38-40] or [17].

Appendix A.2. T matrices

The Hamiltonian matrix in Eq. (12) is constructed by plugging the MT part of Eq. (8) in Eq. (10), and carrying on the integration and summations over a. One first splits the Kohn-Sham Hamiltonian $\hat{H}_{\mathrm{KS}}$ in spherical part H_{sph} (kinetic energy and atomic Coulomb interaction) — which depends only on the radial distance r_a from the center of each MT — and a non-spherical part $V(\mathbf{r})$, involving the rest of the potential which is further expanded in terms of spherical harmonics

$$V(\mathbf{r}) = \sum_{l,m} v_l(r) Y_{l,m}(\hat{r}). \tag{A.3}$$

The net result is the unwieldy expression

$$\begin{aligned}(H)_{t',t} = \sum_a \sum_{L'} \sum_{L} \int \left[A_{L',a,t} u_{l',a}(r) + B_{L',a,t} \dot{u}_{l',a}(r)\right]^* Y^*_{L'}(\hat{r}) \cdot \\ \cdot \left[H_{\mathrm{sph}}(r) + \sum_{L''} v_{l''}(r) Y_{L''}(\hat{r})\right] \cdot \left[A_{L,a,t} u_{l,a}(r) + B_{L,at} \dot{u}_{l,a}(r)\right] Y_L(\hat{r}) \, \mathrm{d}^3 r.\end{aligned} \tag{A.4}$$

By construction, the functions $u_{l,a}$ are the exact solution of the spherical Hamiltonian (see Eq. (7)) so that the elements of H_{sph} contribute solely to the diagonal of the $(H)_{t',t}$ matrix.

The remaining four terms of Eq. (A.4), coming from the non-spherical contribution, fill up densely the rest of the entries of $(H)_{t',t}$. In order to have a more manageble expression, we switch to spherical coordinates, so that we collect the angular dependence on spherical harmonics in a self-contained integral form separated from the radial part

$$G_{L,L',L''} \equiv \int Y^*_{L'} Y_L Y_{L''} \mathrm{d}\Omega, \tag{A.5}$$

where $\mathrm{d}\Omega \equiv \sin\theta \mathrm{d}\phi \mathrm{d}\theta$. Such integrals are called Gaunt coefficients and have a number of symmetries which simplify the numerical generation.

The A and B coefficients do not depend on the radial coordinates and can be factored out. The rest of the functions are then collected in a set of four integrals differing for the contributions from the radial functions u_l and their energy derivatives $\dot{u}_l$

$$\begin{aligned}I^{(u,u)}_{l,l',l'';a} &= \int u_{l',a}(r)\ v_{l''}(r)\ u_{l,s}(r)\ r^2 \mathrm{d}r, \\ I^{(\dot{u},\dot{u})}_{l,l',l'';a} &= \int \dot{u}_{l',a}(r)\ v_{l''}(r)\ \dot{u}_{l,s}(r)\ r^2 \mathrm{d}r, \\ I^{(u,\dot{u})}_{l,l',l'';a} &= \int u_{l',a}(r)\ v_{l''}(r)\ \dot{u}_{l,s}(r)\ r^2 \mathrm{d}r, \\ I^{(\dot{u},u)}_{l,l',l'';a} &= \int \dot{u}_{l',a}(r)\ v_{l''}(r)\ u_{l,s}(r)\ r^2 \mathrm{d}r.\end{aligned} \tag{A.6}$$

By bringing together all the previous mathematical objects — spherical and non-spherical

contribution alike — one defines the set of four T-matrices which appear in Eq. (12)

$$
\begin{aligned}
T^{[AA]}_{L',L;a} &= \sum_{L''} I^{(u,u)}_{L,L',L''} \cdot G_{L,L',L''} \;+\; \delta_{L',L} E_l, \\
T^{[BB]}_{L',L;a} &= \sum_{L''} I^{(\dot{u},\dot{u})}_{L,L',L''} \cdot G_{L,L',L''} \;+\; \delta_{L',L} E_l \left\| \dot{u}_{l;a} \right\|, \\
T^{[AB]}_{L',L;a} &= \sum_{L''} I^{(u,\dot{u})}_{L,L',L''} \cdot G_{L,L',L''} \;+\; \delta_{L',L}, \\
T^{[BA]}_{L',L;a} &= \sum_{L''} I^{(\dot{u},u)}_{L,L',L''} \cdot G_{L,L',L''} \;+\; 0.
\end{aligned}
\tag{A.7}
$$

Acknowledgements

Financial support from the Jülich Aachen Research Alliance-High Performance Computing, the Deutsche Forschungsgemeinschaft (DFG) through grant GSC 111, and from the Deutsche Telekom Stiftung are gratefully acknowledged. We wish to thank Cuauhtemoc Salazar for his thorough review of the manuscript. Furthermore, we thank our colleagues in the HPAC research group for many fruitful discussions and valuable feedback.

References

References

[1] Fleur: The Jülich FLAPW code family (Feb. 2016).
URL `http://www.flapw.de/pm/index.php?n=Main.AboutFleur`

[2] E. Peise, D. Fabregat-Traver, P. Bientinesi, On the Performance Prediction of BLAS-based Tensor Contractions, High Performance Computing Systems. Performance Modeling, Benchmarking, and Simulation 8966 (2015) 193–212.

[3] C. Bischof, C. V. Loan, The wy representation for products of householder matrices, SIAM Journal on Scientific and Statistical Computing 8 (1) (1987) s2–s13. `doi:10.1137/0908009`.
URL `http://dx.doi.org/10.1137/0908009`

[4] T. Joffrain, T. M. Low, E. S. Quintana-Ortí, R. v. d. Geijn, F. G. V. Zee, Accumulating householder transformations, revisited, ACM Trans. Math. Softw. 32 (2) (2006) 169–179. `doi:10.1145/1141885.1141886`.
URL `http://doi.acm.org/10.1145/1141885.1141886`

[5] A. Canning, W. Mannstadt, A. J. Freeman, Parallelization of the FLAPW method, Computer Physics Communications 130 (3) (2000) 233–243.

[6] M. Petersen, F. Wagner, L. Hufnagel, M. Scheffler, P. Blaha, K. Schwarz, Improving the efficiency of FP-LAPW calculations, Computer Physics Communications 126 (3) (2000) 294–309.

[7] F. Nogueira, M. A. L. Marques, C. Fiolhais, A primer in density functional theory, Lecture Notes in Physics, Springer, Berlin, 2003.

[8] D. Sholl, J. A. Steckel, Density Functional Theory, A Practical Introduction, John Wiley & Sons, 2011.

[9] E. Wimmer, H. Krakauer, M. Weinert, A. J. Freeman, Full-Potential Self-Consistent Linearized-Augmented-Plane-Wave Method for Calculating the Electronic-Structure of Molecules and Surfaces - O2 Molecule, Physical Review B 24 (2) (1981) 864–875.

[10] H. J. F. Jansen, A. J. Freeman, Total-Energy Full-Potential Linearized Augmented-Plane-Wave Method for Bulk Solids - Electronic and Structural-Properties of Tungsten, Physical Review B 30 (2) (1984) 561–569.

[11] K. Burke, Perspective on density functional theory, The Journal of Chemical Physics 136 (15) (2012) 150901.

[12] P. Hohenberg, Inhomogeneous Electron Gas, Physical Review 136 (3B) (1964) B864–B871.

[13] W. Kohn, L. J. Sham, Self-Consistent Equations Including Exchange and Correlation Effects, Phys.Rev. 140 (1965) A1133–A1138.

[14] M. Weinert, Solution of Poisson's equation: Beyond Ewald–type methods, Journal of Mathematical Physics.

[15] N. Ashcroft, N. Mermin, Solid State Physics, HRW international editions, Holt, Rinehart and Winston, 1976.

[16] P. Kurz, Non-collinear magnetism at surfaces and in ultrathin films, Ph.D. thesis, RWTH Aachen (Feb. 2000).
URL `http://www.fz-juelich.de/pgi/pgi-1/DE/Leistungen/MasterDiplomDr/_node.html`

[17] D. J. Singh, L. Nordström, Planewaves, Pseudopotentials and the LAPW Method, Springer US, 2006.

[18] E. Di Napoli, E. Peise, M. Hrywniak, P. Bientinesi, High-performance generation of the hamiltonian and overlap matrices in flapw methods, arXiv:1602.06589 [cs, physics].
URL `https://arxiv.org/abs/1602.06589`

[19] E. Di Napoli, M. Berljafa, Block iterative eigensolvers for sequences of correlated eigenvalue problems, Computer Physics Communications 184 (11) (2013) 2478–2488.

[20] M. Berljafa, D. Wortmann, E. Di Napoli, An optimized and scalable eigensolver for sequences of eigenvalue problems, Concurrency and Computation: Practice and Experience 27 (4) (2014) 905–922.

[21] P. Blaha, H. Hofstätter, O. Koch, R. Laskowski, K. Schwarz, Iterative diagonalization in augmented plane wave based methods in electronic structure calculations, Journal of Computational Physics 229 (2) (2010) 453–460.

[22] D. E. Knuth, Structured programming with go to statements, ACM Comput. Surv. 6 (4) (1974) 261–301.

[23] J. Poulson, B. Marker, R. A. van de Geijn, J. R. Hammond, N. A. Romero, Elemental: A new framework for distributed memory dense matrix computations, ACM Trans. Math. Softw. 39 (2) (2013) 13:1–13:24. doi:10.1145/2427023.2427030.
URL http://doi.acm.org/10.1145/2427023.2427030